THE LOCAL MAXIMA OF MAXIMAL INJECTIVITY RADIUS AMONG HYPERBOLIC SURFACES

JASON DEBLOIS

ABSTRACT. The function on the Teichmüller space of complete, orientable, finite-area hyperbolic surfaces of a fixed topological type that assigns to a hyperbolic surface its maximal injectivity radius has no local maxima that are not global maxima.

Let $\mathfrak{T}_{g,n}$ be the Teichmüller space of complete, orientable, finite-area hyperbolic surfaces of genus g with n cusps. In this paper we begin to analyze the function $max\colon \mathfrak{T}_{g,n} \to \mathbb{R}^+$ that assigns to $S \in \mathfrak{T}_{g,n}$ its maximal injectivity radius. The *injectivity radius* of S at x, $injrad_x(S)$, is half the length of the shortest non-constant geodesic arc in S with both endpoints at x. It is not hard to see that $injrad_x(S)$ varies continuously with x and approaches 0 in the cusps of S, so it attains a maximum on any fixed finite-area hyperbolic surface S.

Our main theorem characterizes local maxima of max on $\mathfrak{T}_{g,n}$:

Theorem 0.1. *For $S \in \mathfrak{T}_{g,n}$, the function max attains a local maximum at S if and only if for each $x \in S$ such that $injrad_x(S) = max(S)$, each edge of the Delaunay tessellation of (S,x) has length $2injrad_x(S)$ and each face is a triangle or monogon.*

Here for a hyperbolic surface S with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to S$, and $x \in S$, the *Delaunay tessellation of (S,x)* is the projection to S of the Delaunay tessellation of $\pi^{-1}(x) \subset \mathbb{H}^2$, as defined by an empty circumcircles condition (see Section 2 below). In particular, a *monogon* is the projection to S of the convex hull of a P-orbit in $\pi^{-1}(x)$, for a maximal parabolic subgroup P of $\pi_1 S$ acting on $\mathbb{H}^2$ by covering transformations.

Theorem 5.11 of the author's previous paper [3] characterized the global maxima of max by a condition equivalent to that of Theorem 0.1, extending work of Bavard [1]. We thus have:

Corollary 0.2. *All local maxima of max on $\mathfrak{T}_{g,n}$ are global maxima.*

This contrasts the behavior of $syst$, the function on $\mathfrak{T}_{g,n}$ that records the *systole*, ie. shortest geodesic, length of hyperbolic surfaces: P. Schmutz Schaller proved in [10] that for many g and n, $syst$ has local maxima on $\mathfrak{T}_{g,n}$ that are not global maxima. Comparing with $syst$, which is well-studied, is one motivation for studying max. (Note that for a closed hyperbolic surface S, $syst(S)$ is twice the *minimal* injectivity radius of S.)

The referee has sketched a direct argument to show that max attains a global maximum on $\mathfrak{T}_{g,n}$. (This is also sketched in the preprint [7], and I prove a somewhat more general fact as Proposition 4.3 of [5].) Together with this observation, Theorem 0.1 gives an alternative proof of Theorem 5.11 of [3], which is not completely independent of the results of [3] but uses only some early results from Sections 1 and Section 2.1 there.

We prove Theorem 0.1 by describing explicit, injectivity radius-increasing deformations of pointed surfaces (S,x) that do not satisfy its criterion. The deformations are produced

by changing finite edge lengths of a decomposition $\mathcal{T}$ of S into compact and horocyclic ideal triangles, with vertex set x. In Section 1 we introduce a space $\mathfrak{D}(S, \mathcal{T})$ parametrizing such deformations. Proposition 1.3 there shows that the natural map $\mathfrak{D}(S, \mathcal{T}) \to \mathfrak{T}_{g,n}$ is continuous. Proposition 1.4 gives a simple description of max near $\mathbf{d} \in \mathfrak{D}(S, \mathcal{T})$ in terms of the edge lengths, assuming that all shortest arcs based at x are edges of $\mathcal{T}$.

By Lemma 2.3, all such arcs are Delaunay edges. Section 2 introduces the Delaunay tessellation of (S, x), following the author's prior paper [2], and describes its relevance to this paper. In particular, we prove Theorem 0.1 using triangulations obtained by subdividing the Delaunay tessellation, see Lemma 2.4.

Section 3 is devoted to the proof of Theorem 0.1. Proposition 3.1 reduces it to the case that every Delaunay edge has length $2injrad_x(S)$. We believe it has more to say about the critical set of max (properly interpreted, since max is not smooth) and hope in future work to more deeply understand this set. Theorem 0.1 is then proved by showing that complicated Delaunay cells can be broken apart by injectivity radius-increasing deformations. The arguments in this section use some basic observations from [3] and, like the results there, exploit what you might call the "calculus of cyclic polygons" laid out in [4].

Acknowledgements. We were originally motivated and in part inspired by a private communication from Ian Agol, where he sketched a proof of [3, Theorem 5.11] using deformations through hyperbolic cone surfaces and conjectured Corollary 0.2. We thank Ian for his interest and ideas. We are also grateful to the referee for helpful comments which have significantly improved the paper, in particular for the simplified proof of Proposition 1.3.

After this paper was first submitted, M. Gendulphe posted the preprint [7] which proves Theorem 0.1 by a different method.

1. Deformations via triangulations

Let us begin this section by fixing a complete, oriented, finite-area hyperbolic surface S of genus g with n cusps and a decomposition $\mathcal{T}$ of S into compact and horocyclic ideal hyperbolic triangles. (Here a *horocyclic ideal triangle* has two vertices on a horocycle of $\mathbb{H}^2$ and an ideal vertex at the horocycle's ideal point.) We will call the pair $(S, \mathcal{T})$ a *triangulated surface* for short. The main results of this section are Propositions 1.3 and 1.4.

We first define a space $\mathfrak{D}(S, \mathcal{T})$ of possible deformations of the edge lengths of $\mathcal{T}$.

Definition 1.1. Suppose $(S, \mathcal{T})$ is a complete, oriented, triangulated hyperbolic surface of finite area, and enumerate the faces of $\mathcal{T}$ as $F_1, \ldots, F_k$ and the compact edges as $e_1, \ldots, e_l$. Let $U \subset (0, \infty)^l$ be the set of $\mathbf{d} = (d_0, \ldots, d_l)$ such that $d_{j_1} < d_{j_2} + d_{j_3}$ for any j_1, j_2, j_3 such that e_{j_1}, e_{j_2} and e_{j_3} are the distinct edges of some F_i. For each vertex x of $\mathcal{T}$ define:

$$A_x(\mathbf{d}) = \sum_i \cos^{-1}\left(\frac{\cosh d_{j_1} \cosh d_{j_2} - \cosh d_{j_3}}{\sinh d_{j_1} \sinh d_{j_2}}\right) + \sum_k \sin^{-1}\left(\frac{1}{\cosh(d_{j_k}/2)}\right)$$

This sum is taken over all i such that x is in the compact triangle F_i, where e_{j_1} and e_{j_2} are the edges of $\mathcal{T}$ containing x, and e_{j_3} is the edge of F_i opposite x, and all k such that x is in the horocyclic triangle F_k with finite side length d_{j_k}. We then take:

$$\mathfrak{D}(S, \mathcal{T}) = \{\mathbf{d} \in U \mid A_x(\mathbf{d}) = 2\pi \text{ for each vertex } x \text{ of } \mathcal{T}\}$$

We call this the set of *deformations of* $(S, \mathcal{T})$.

Lemma 1.2. *Suppose $(S, \mathcal{T})$ is a complete, oriented, triangulated hyperbolic surface of finite area. For each face F_i of $\mathcal{T}$ and each $\mathbf{d} = (d_1, \ldots, d_l) \in \mathfrak{D}(S, \mathcal{T})$ let $F_i(\mathbf{d})$ be the compact hyperbolic triangle with edge lengths d_{j_1}, d_{j_2} and d_{j_3}, if F_i is compact with edges e_{j_1}, e_{j_2} and e_{j_3}; or otherwise the horocyclic ideal triangle with finite edge length d_{j_1}, where F_i has compact edge e_{j_1}. The triangulated polyhedral complex $(S(\mathbf{d}), \mathcal{T}(\mathbf{d}))$ obtained by identifying edges of the $F_i(\mathbf{d})$ in pairs corresponding to edges of $\mathcal{T}$ inherits a complete hyperbolic structure from the $F_i(\mathbf{d})$, and it has a homeomorphism to $(S, \mathcal{T})$ taking $F_i(\mathbf{d})$ to F_i for each i.*

Proof. Let us recall some standard facts. Below, for a compact hyperbolic triangle with sides of length a, b and c let α be the interior angle opposite the side of length a. Let δ be the interior angle at either endpoint of the finite edge, with length d, of a horocyclic ideal triangle. Then:

$$(1.2.1) \quad \alpha = \cos^{-1}\left(\frac{\cosh b \cosh c - \cosh a}{\sinh b \sinh c}\right) \in (0, \pi) \quad \delta = \sin^{-1}\left(\frac{1}{\cosh(d/2)}\right) \in (0, \pi/2)$$

The left-hand equation is the hyperbolic law of cosines (see eg. [9, Theorem 3.5.3]). The right can be proved by an explicit calculation in, say, the upper half-plane model $\mathbb{R} \times (0, \infty)$ for $\mathbb{H}^2$, placing the horocycle at $\mathbb{R} \times \{1\}$ and using the fact that the Euclidean and hyperbolic distances ℓ and d, respectively, between points on it satisfy $\ell/2 = \sinh(d/2)$.

Now with the faces and compact edges of $\mathcal{T}$ enumerated as in Definition 1.1, for $\mathbf{d} \in \mathfrak{D}(S, \mathcal{T})$ and $1 \leq i \leq k$ let $F_i(\mathbf{d})$ be as described in the statement. By construction and the formulas of (1.2.1), $A_x(\mathbf{d}) = 2\pi$ is the sum of all vertex angles of the $F_i(\mathbf{d})$. Note also that each edge of $\mathcal{T}$ is contained in two faces, and again by construction if this edge is compact then the corresponding edges of the $F_i(\mathbf{d})$ have the same length.

For $(S(\mathbf{d}), \mathcal{T}(\mathbf{d}))$ as described in the statement, there is clearly a triangulation-preserving homeomorphism $(S(\mathbf{d}), \mathcal{T}(\mathbf{d})) \to (S, \mathcal{T})$. Moreover, choosing a disjoint collection of representatives of the $F_i(\mathbf{d})$ in $\mathbb{H}^2$ it is not hard to arrange for the pairing of edges to be realized by an $\mathrm{Isom}^+(\mathbb{H}^2)$-*side pairing* in the sense of [9, §9.2]. Theorem 9.2.2 of [9] then implies that $S(\mathbf{d})$ inherits a hyperbolic structure from the $F_i(\mathbf{d})$. The key requirement for this result, that the side-pairing is *proper*, obtains from the fact that the angle sum $A_x(\mathbf{d})$ around the vertex x is 2π.

We further claim that the hyperbolic structure on $S(\mathbf{d})$ is complete, see [9, Theorem 8.5.9]. This follows from the stipulation in Definition 1.1 that non-compact faces of $\mathcal{T}$ are horocyclic ideal triangles, since an isometry that takes an infinite edge of one horocyclic ideal triangle to an infinite edge of another identifies the horocycles containing their vertices. For any such face with ideal vertex v, we thus have $d(v) = 0$, where $d(v)$ is the "gluing invariant" of [9, §9.8], so S is complete by Theorem 9.8.5 of [9] (cf. [11, Prop. 3.4.18]). $\square$

We next relate the deformation space $\mathfrak{D}(S, \mathcal{T})$ to the Teichmüller space $\mathfrak{T}_{g,n}$ of hyperbolic surfaces with the topological type of S, endowed with its standard topology (see eg. [6]). Here we will regard the hyperbolic surfaces $S(\mathbf{d})$ from Lemma 1.2 as marked by the homeomorphism from S described there.

Proposition 1.3. *For a complete, oriented, triangulated hyperbolic surface $(S, \mathcal{T})$ of finite area, with genus g and n cusps, the map $\mathfrak{D}(S, \mathcal{T}) \to \mathfrak{T}_{g,n}$ given by $\mathbf{d} \mapsto S(\mathbf{d})$ is continuous.*

Proof. We will show that for any essential simple closed curve γ on S, the function $\mathbf{d} \mapsto \ell_{S(\mathbf{d})}(\gamma)$ that measures the geodesic length of γ in $S(\mathbf{d})$ is continuous. It then follows from standard results, eg. the "$9g - 9$ theorem" [6, Theorem 10.7], that $\mathbf{d} \mapsto S(\mathbf{d})$ is continuous.

Fix $\mathbf{d} \in \mathfrak{D}(S, \mathcal{T})$, and refer by γ to an oriented geodesic representative of γ on $S(\mathbf{d})$. For $\mathbf{d}'$ near to $\mathbf{d}$, we now construct a piecewise-geodesic γ' on $S(\mathbf{d}')$ which will be evidently isotopic to the image of γ under the homeomorphism $S(\mathbf{d}) \to S(\mathbf{d}')$ described in Lemma 1.2. We will then show that the length of γ' exceeds that of γ by no more than some ϵ depending on $\delta = \max\{|\delta_i|\}$, where $\delta_i = |d_i - d'_i|$ for each i, which approaches 0 as $\delta \to 0$.

Partition γ into arcs $\gamma_0, \ldots, \gamma_{k-1}$ with disjoint interiors such that γ_j is adjacent to γ_{j+1} for each j, and each γ_j is either an edge of or (the generic case) properly embedded in a triangle of $\mathcal{T}$. For each j we construct a geodesic arc γ'_j in $S(\mathbf{d}')$ that lies in the same triangle(s) of $\mathcal{T}$ as γ_j, as follows. For each endpoint x of γ_j that is a vertex of $\mathcal{T}$, let the corresponding endpoint x' of γ'_j be the same vertex; if x lies in the interior of a compact edge e_{i_j} of $\mathcal{T}$, then with δ_{i_j} as above let x' lie on e_{i_j} in $S(\mathbf{d}')$ with $|d(x', v) - d(x, v)| = \delta_{i_j}/2$ for each vertex v of e_{i_j}; and if x is in the interior of a non-compact edge of $\mathcal{T}$ then let x' lie on the same edge in $S(\mathbf{d}')$, at the same distance from its (finite) vertex. Now let $\gamma' = \gamma'_1 \cup \ldots \gamma'_{k-1}$.

For any fixed j, we will show that $|\ell'_j - \ell_j| \to 0$ as $\delta \to 0$, where ℓ_j and ℓ'_j are the respective lengths of γ_j and γ'_j in $S(\mathbf{d})$ and $S(\mathbf{d}')$. If γ_j is an edge of $\mathcal{T}$ this is obvious, so let us assume it is not. Then γ_j cuts the triangle T of $\mathcal{T}$ containing it into two pieces, at least one of which is a triangle. If the vertex v of the sub-triangle of T that does not lie in γ is finite, and the interior angle there is θ_j, then by the hyperbolic law of cosines γ_j has length ℓ_j given by

$$\cosh \ell_j = \cosh d(x_j, v) \cosh d(x_{j+1}, v) - \sinh d(x_j, v) \sinh d(x_{j+1}, v) \cos \theta_j.$$

Here x_j and x_{j+1} are the endpoints of γ_j, and θ_j is given in terms of $\mathbf{d}$ by the left side of the formula (1.2.1). For ℓ'_j we substitute x'_j, x'_{j+1} and θ'_j above. It is clear from this formula that $|\ell'_j - \ell_j| \to 0$ as $\delta \to 0$.

If the vertex v described above is ideal then we claim that the length ℓ_j of γ_j satisfies

$$\cosh \ell_j = \frac{e^{-d(x_j, v_j) - d(x_{j+1}, v_{j+1})}}{2} \left(4 \sinh^2(d_{i_j}/2) + e^{2d(x_j, v_j)} + e^{2d(x_{j+1}, v_{j+1})}\right),$$

where v_j is the finite vertex of the edge e of T containing x_j, v_{j+1} is the finite vertex of the edge containing x_{j+1}, and d_{i_j} is the length of the compact edge of T.

This follows from explicit computations in the upper half-plane model. Applying an isometry, we may take T inscribed in the horocycle $C = \mathbb{R} + i$, with $v_j = i$ and $v_{j+1} = \theta + i$, where $\theta = 2 \sinh(d_{i_j}/2)$ is the distance from v_j to v_{j+1} along C. Then $x_j = iy_0$ and $x_{j+1} = \theta + iy_1$ for $y_0, y_1 > 1$ satisfying $e^{d(x_j, v_j)} = y_0$ and $e^{d(x_{j+1}, v_{j+1})} = y_1$. Theorem 1.2.6(ii) of [8] now proves the claim, giving:

$$\cosh \ell = 1 + \frac{\theta^2 + (y_0 - y_1)^2}{2y_0 y_1} = \frac{\theta^2 + y_0^2 + y_1^2}{2y_0 y_1}$$

To compute the length ℓ'_j of the corresponding arc γ'_j we simply replace d_{i_j} by d'_{i_j} above, where d'_{i_j} is the length in $S(\mathbf{d}')$ of the compact edge of T. Convergence of ℓ'_j to ℓ_j thus follows as in the previous case.

Since the length of γ is $\sum \ell_i$, and the geodesic length of γ' is at most $\sum \ell'_i$, this exceeds the length of γ by no more than some $\epsilon = \epsilon(\delta)$, which approaches 0 as $\delta \to 0$. From the formulas

above we see that the dependence of ϵ on δ is uniform on compact subsets of the open set U of Definition 1.1. Therefore we can apply the same argument with the roles of γ and γ' reversed to show that the length of γ exceeds the geodesic length of γ' by no more than some $\epsilon' = \epsilon(\delta)$ which also approaches 0 as $\delta \to 0$. Continuity follows. $\qquad\square$

Proposition 1.4. *Suppose $(S, \mathcal{T})$ is a complete, triangulated hyperbolic surface of finite area with vertex set $\{x\}$ such that the entire collection of geodesic arcs of length $2\,injrad_x(S)$ based at x is a set $e_{j_1}, \ldots, e_{j_n}$ of edges of $\mathcal{T}$. Then there is a neighborhood V in $\mathfrak{D}(S, \mathcal{T})$ of the edge length collection of $\mathcal{T}$ such that for any $\mathbf{d} \in V$, if $x_{\mathbf{d}}$ is the vertex of the triangulated hyperbolic surface $(S(\mathbf{d}), \mathcal{T}(\mathbf{d}))$ of Lemma 1.2 then $injrad_{x(\mathbf{d})} S(\mathbf{d}) = \frac{1}{2} \min\{d_{j_i}\}_{i=1}^n$.*

Proof. The collection of geodesic arcs in S based at x is in 1-1 correspondence with those in $\mathbb{H}^2$ joining $\tilde{x}$ to other points of $\pi^{-1}(x)$, where $\pi \colon \mathbb{H}^2 \to S$ is a locally isometric universal cover and $\tilde{x}$ is a fixed element of $\pi^{-1}(x)$. Fix some $R > 0$ that is slightly larger than $2\,injrad_x(S)$, and let $P \subset \mathbb{H}^2$ be the union of lifts of triangles of $\mathcal{T}$ that intersect the closed ball of radius R about $\tilde{x}$. This is a finite union since the lifted triangulation is locally finite.

For $\mathbf{d}'$ near the edge length collection $\mathbf{d}$ of S in $\mathfrak{D}(S, \mathcal{T})$, let P' be the corresponding union of triangles lifted from $(S(\mathbf{d}'), \mathcal{T}(\mathbf{d}'))$. That is, fix a locally isometric universal cover $\pi' \colon \mathbb{H}^2 \to S(\mathbf{d}')$ and some $\tilde{x}' \in (\pi')^{-1}(x')$, where x' is the vertex of $\mathcal{T}(\mathbf{d}')$, and let P' be the image of P under the lift that takes $\tilde{x}$ to $\tilde{x}'$ of the marking $S \to S(\mathbf{d}')$. (Recall that this map takes triangles to triangles.)

For each geodesic arc γ in S based at x and any fixed $\epsilon > 0$, arguing as in the proof of Proposition 1.3 shows that $\mathbf{d}'$ can be chosen near enough to $\mathbf{d}$ that the geodesic arc in $S(\mathbf{d}')$ based at x' in the based homotopy class of the image of γ has length less than ϵ away from the length of γ. In particular, there exists $\delta > 0$ such that if $\max\{|d_i - d_i'|\} < \delta$ then for each vertex v of P at distance greater than $2\,injrad_x(S)$ from $\tilde{x}$, the image of v has distance to $\tilde{x}'$ greater than $i(\mathbf{d}') \doteq \min\{\ell_{j_1}', \ldots, \ell_{j_n}'\}$, where $\ell_{j_i'}$ is the length of e_{j_i} in $S(\mathbf{d}')$ for each i.

We now consider geodesic arcs in S based at x whose lifts based at $\tilde{x}$ exit P. For each such arc the analogous fact holds for its correspondent in $S(\mathbf{d}')$. We will thus complete the proposition's proof by showing that for small enough $\delta > 0$, the closest point to $\tilde{x}'$ on each edge in the frontier of P' is at distance greater than $i(\mathbf{d}')$ (defined above) from it, whence $i(\mathbf{d}')$ is twice the injectivity radius of $S(\mathbf{d}')$ at x'.

For a compact edge e in the frontier of P, let T be the triangle with one edge at e and opposite vertex $\tilde{x}$. The closest point of e to $\tilde{x}$ is in its interior if and only if the angles of T are less than $\pi/2$ at each endpoint of e. In this case the geodesic arc from $\tilde{x}$ to its closest point on e intersects e at right angles, and by the hyperbolic law of sines the distance h from $\tilde{x}$ to e satisfies $\sinh h = \sinh \ell \sin \theta$. Here ℓ is the distance from $\tilde{x}$ to an endpoint v of e, and θ is the interior angle of T at v.

If T' is the corresponding triangle in P' then for $\mathbf{d}'$ near $\mathbf{d}$, each edge length of T' is near the corresponding edge length of T, as we have already remarked, and it follows from the hyperbolic law of cosines that the same holds for the angles of T and T'. In particular, if the closest point of e to $\tilde{x}$ is in the interior of e then for $\mathbf{d}'$ near enough to $\mathbf{d}$, the closest point of the corresponding edge e' to $\tilde{x}'$ is also in its interior, and by the hyperbolic law of sines the distance h' from $\tilde{x}'$ to e' approaches h as $\mathbf{d}' \to \mathbf{d}$. The remaining case is straightforward.

For a non-compact edge e in the frontier of P, if the nearest point of e to $\tilde{x}$ is in its interior then we again use the formula $\sinh h = \sinh \ell \sin \phi$, where now ℓ is the length of the geodesic

arc from $\tilde{x}$ to the finite endpoint v of e, and ϕ is the angle from this arc to e. It follows as before that $\ell' \to \ell$ as $\mathbf{d}' \to \mathbf{d}$. To see that the corresponding angle ϕ' approaches to ϕ as $\mathbf{d}' \to \mathbf{d}$ we note that $\phi = \theta + \delta$, where δ is the interior angle at v of the horocyclic triangle T_0 in P containing e, and θ is the interior angle at v of the triangle determined by $\tilde{x}$ and the finite side f of T_0. The corresponding angles $\delta' \to \delta$ and $\theta' \to \theta$, so $\phi' \to \phi$ as $\mathbf{d}' \to \mathbf{d}$. $\qquad\square$

2. The Delaunay tessellation

In this section, for a hyperbolic surface S and $x \in S$ we define the Delaunay tessellation of (S, x) (Definition 2.2 below) by projecting Delaunay cells of $\pi^{-1}(x)$, where $\pi\colon \mathbb{H}^2 \to S$ is the universal cover. Here the Delaunay tessellation of a locally finite, lattice-invariant subset $\widetilde{\mathcal{S}} \subset \mathbb{H}^2$, in the sense of [3, Theorem 5.1] (which itself is the specialization to two dimensions of [2, Theorem 6.23]), is characterized by the *empty circumcircles condition*:

> For each circle or horocycle H of $\mathbb{H}^2$ that intersects $\widetilde{\mathcal{S}}$ and bounds a disk or horoball B with $B \cap \widetilde{\mathcal{S}} = H \cap \widetilde{\mathcal{S}}$, the closed convex hull of $H \cap \widetilde{\mathcal{S}}$ in $\mathbb{H}^2$ is a Delaunay cell. Each Delaunay cell has this form.

In proving Theorem 0.1 we will use triangulations *compatible* with the Delaunay tessellation of (S, x), in the sense of Lemma 2.4. There are three advantages to working with the Delaunay tessellation. First, every geodesic arc of length $2injrad_x(S)$ based at x is a Delaunay edge, as we prove in Lemma 2.3. Second, by construction Delaunay cells are cyclic or horocyclic; that is, inscribed in metric circles or horocycles, respectively. In [4] there are calculus formulas describing the derivative of area with respect to side length for such polygons.

Finally, the Delaunay tessellation of $\mathcal{S}$ contains the geometric dual to the *Voronoi tessellation* of $\mathcal{S}$, which has two-cells of the form

$$V_{\mathbf{x}} = \{y \in \mathbb{H}^2 \mid d(y, \mathbf{x}) \le d(y, \mathbf{x}') \ \forall \ \mathbf{x}' \in \mathcal{S}\},$$

for each $\mathbf{x} \in \mathcal{S}$. See eg. [2, §5]. Its edges are intersections $V_{\mathbf{x}} \cap V_{\mathbf{x}'}$ containing at least two points. The *geometric dual* to any such edge is the geodesic arc joining $\mathbf{x}$ to $\mathbf{x}'$. In Section 3 we will exploit the geometric duality relation using some results from [3, §2.1] that show how the Voronoi tessellation encodes certain extra structure associated to "non-centered" Delaunay two-cells. This helps us overcome the central difficulty in using deformations via triangulations, which is that the area of cyclic polygons is not monotonic in their side lengths.

The first result we will prove here is mostly [3, Corollary 5.2], which is again the specialization of a result from [2], Corollary 6.27 there. Theorem 5.1 of [3] asserts for a set $\widetilde{\mathcal{S}}$ invariant under a lattice Γ that a Delaunay cell of $\widetilde{\mathcal{S}}$ is inscribed in a horocycle C if and only if its stabilizer in Γ is a parabolic subgroup Γ_u of Γ that fixes the ideal point C. Such cells are the primary concern of this result.

Corollary 2.1. *For a complete, oriented, finite-area hyperbolic surface F with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$, and a finite set $\mathcal{S} \subset F$, there are finitely many $\pi_1 F$-orbits of Delaunay cells of $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$. The interior of each compact Delaunay cell embeds in F under π. For a cell C_u with parabolic stabilizer Γ_u, $\pi|_{int\,C_u}$ factors through an embedding of $int\,C_u/\Gamma_u$ to a set containing a cusp of F.*

A fundamental domain in a parabolic-invariant cell C_u for the action of its stabilizer Γ_u is a horocyclic ideal polygon whose finite-length edges are edges of C_u.

Proof. The first part of this result was recorded as Corollary 5.2 of [3]. The second part follows from Lemma 5.7 there. This result implies that the vertices of C_u can be enumerated as $\{s_i \,|\, i \in \mathbb{Z}\}$ so that s_i and s_{i+1} bound an edge γ_i of C_u for each i, and $g(s_i) = s_{i+k}$ for some fixed $k \in \mathbb{Z}$, where g is the generator of Γ_u. It follows that a fundamental domain for the Γ_u-action is the non-overlapping union of horocyclic triangles $T_i \cup T_{i+1} \ldots \cup T_{i+k-1}$ defined in Lemma 5.7 for any fixed i. This is a horocyclic ideal $(k+1)$-gon, see [4, Prop. 3.8]. $\qquad\square$

Definition 2.2. For a complete, oriented, finite-area hyperbolic surface S and $x \in S$, we will call the *Delaunay tessellation of* (S, x) the projection to S of the Delaunay tessellation of $\pi^{-1}(x)$, for some fixed universal cover $\pi \colon \mathbb{H}^2 \to S$.

Lemma 2.3. *For a complete, oriented, finite-area hyperbolic surface S and $x \in S$, every geodesic arc based at x with length $2injrad_x(S)$ is an edge of the Delaunay tessellation of (S, x). In particular, the injectivity radius of S at x is half the minimum edge length of the Delaunay tessellation.*

Proof. Every Delaunay edge of (S, x) is a non-constant geodesic arc with both endpoints at x, so its length is at least $2injrad_x(S)$. For a closed geodesic arc γ of length $2injrad_x(S)$ based at x, let $\tilde{\gamma}$ be a lift of γ to $\mathbb{H}^2$. The metric disk B of radius $injrad_x(S)$ centered at the midpoint of $\tilde{\gamma}$ intersects $\pi^{-1}(x)$ in the endpoints of $\tilde{\gamma}$. Every other point of B has distance less than $2injrad_x(S)$ from the endpoints of $\tilde{\gamma}$, so $B \cap \pi^{-1}(x) = \partial\tilde{\gamma}$. It follows that $\tilde{\gamma}$ is a Delaunay edge of $\pi^{-1}(x)$, hence that γ is a Delaunay edge of (S, x). $\qquad\square$

Lemma 2.4. *For any complete, oriented, hyperbolic surface S of finite area and $x \in S$, there is a decomposition $\mathcal{T}$ of S into compact and horocyclic ideal triangles that is compatible with the Delaunay tessellation of (S, x) in the sense that its vertex set is $\{x\}$ and each edge of the Delaunay tessellation is an edge of $\mathcal{T}$.*

Proof. $\mathcal{T}$ is compatible with the Delaunay tessellation if its faces are obtained by subdividing Delaunay two-cells into triangles. This can be done for instance by dividing each compact two-cell by diagonals from a fixed vertex, and each horocyclic two-cell into horocyclic ideal triangles. On a horocyclic cell C_u, the latter operation joins each vertex of the fundamental domain for Γ_u of Corollary 2.1 to the ideal point of its circumscribed horocycle. $\qquad\square$

3. Increasing injectivity radius

The goal of this section is to prove the main Theorem 0.1. We will do this in two steps. The first, Proposition 3.1 below, reduces to the case that all compact Delaunay edges have equal length. We then prove the Theorem by addressing the case that all Delaunay edge lengths are equal but there is a complicated Delaunay cell C.

Proposition 3.1. *For a complete, oriented hyperbolic surface S of finite area and $x \in S$ such that $max(S) = injrad_x(S)$, if the Delaunay tessellation of (S, x) has an edge of length greater than $2injrad_x(S)$ then S is not a local maximum of max on $\mathfrak{T}_{g,n}$.*

In fact, there is a continuous map $t \mapsto S_t \in \mathfrak{T}_{g,n}$ on $(-\epsilon, \epsilon)$ for some $\epsilon > 0$, and $x_t \in S_t$ for each t, such that $S_0 = S$, $x_0 = x$, and $\frac{d}{dt} injrad_{x_t}(S_t) = \frac{1}{2}$.

Proof. Let $\mathcal{T}$ be a triangulation compatible with the Delaunay tessellation of (S, x). Enumerate the edges of $\mathcal{T}$ as $\gamma_1, \ldots, \gamma_l$ so that the Delaunay edges consist of those with $j \leq n$ for some $n \leq l$, and γ_j has length $2injrad_x(S)$ if and only if $j \leq m$ for some $m < n$, and let $\mathbf{d} = (d_0, \ldots d_l)$ be the collection of edge lengths. We will produce a smooth map $t \mapsto \mathbf{d}(t) = (d_1(t), \ldots, d_n(t)) \in \mathfrak{D}_D(S, \mathcal{T})$ on some interval $(-\epsilon, \epsilon)$, with $\mathbf{d}(0) = \mathbf{d}$, by prescribing the $d_j(t)$ as follows: take $d_j(t) = d_j + t$ for all $j \leq m$ and leave all other edge lengths $d_j(t)$ constant except for $d_n(t)$, which is determined by the ODE $\frac{d}{dt} A_x(\mathbf{d}(t)) = 0$.

Here A_x is from Definition 1.1. If there is a smooth solution $d_n(t)$ then $A_x(\mathbf{d}(t)) \equiv 2\pi$ since $\mathbf{d}$ is the edge length collection of the triangulated hyperbolic surface $(S, \mathcal{T})$. It will then follow from Proposition 1.3 that $S_t \doteq S(\mathbf{d}(t))$ is a deformation of S in $\mathfrak{T}_{g,n}$, and from Proposition 1.4 that $\frac{d}{dt} injrad_{x_t} S_t = \frac{1}{2}$, where x_t is the vertex of $S(\mathbf{d}(t))$. To show that $\frac{d}{dt} A_x(\mathbf{d}(t)) = 0$ has a smooth solution we rearrange it using the chain rule and our stipulations on the $d_j(t)$, yielding:

$$0 = \sum_{j=1}^{m} \frac{\partial}{\partial d_j} \left(D_0(T_{i_j^+}(\mathbf{d}(t))) + D_0(T_{i_j^-}(\mathbf{d}(t))) \right)$$
$$+ d_n'(t) \cdot \frac{\partial}{\partial d_n} \left(D_0(T_{i_n^+}(\mathbf{d}(t))) + D_0(T_{i_n^-}(\mathbf{d}(t))) \right).$$

Here for each j, $T_{i_j^+}$ and $T_{i_j^-}$ are the triangles containing the edge γ_j; for a triangle T_i with edges γ_{j_1}, γ_{j_2}, γ_{j_3} we refer by $T_i(\mathbf{d}(t))$ to the triple $(d_{j_1}(t), d_{j_2}(t), d_{j_3}(t))$ of changing edge lengths; and $D_0(a, b, c)$ records the area of the triangle with edge lengths a, b and c. In [4] we gave formulas for the partial derivatives of D_0 with respect to a, b and c.

If the coefficient $\frac{\partial}{\partial d_n} \left(D_0(T_{i_n^+}(\mathbf{d}(t))) + D_0(T_{i_n^-}(\mathbf{d}(t))) \right)$ is non-zero then we can solve for $d_n'(t)$, yielding a first-order ODE in $d_n(t)$. We claim this holds at $t = 0$, ie for the $T_{i_n^\pm}(\mathbf{d})$, and therefore at all possible values of $\mathbf{d}(t)$ near $\mathbf{d}$. Given the claim, Picard's theorem on the existence of solutions to first-order ODE implies there is a smooth solution $d_n(t)$ for small t (note that smoothness of D_0 is proven in [4]). We will apply results from Section 2 of [3], together with [4, Proposition 2.3], to prove the claim.

There are two cases, divided by the qualitative nature of the Delaunay cells $C_{i_n}^{\pm}$ of (S, x) containing the triangles $T_{i_n}^{\pm}$. In the first case one of the $C_{i_n^\pm}$, say $C_{i_n^-}$, is compact and therefore cyclic but not centered, and γ_n is its longest side. Here a cyclic polygon is *centered* if its interior contains the center of its circumcircle. The longest side of a non-centered cyclic polygon separates its interior from the center of its circumcircle [4, Prop. 2.2], so since $T_{i_n}^{-}$ is contained in $C_{i_n}^{-}$ it is also non-centered with longest side γ_n.

In this case Lemma 2.5 of [3] asserts that the dual Voronoi vertex v to $C_{i_n^-}$ is the initial vertex of a non-centered Voronoi edge e geometrically dual to γ_n. If e is compact then its terminal vertex is the geometric dual to $C_{i_n^+}$, so by the same result it is not also the case that $C_{i_n^+}$ is non-centered with longest edge γ_{j_0}. This is therefore also not the case for $T_{i_n^-}$, so be [4, Proposition 2.3] the coefficient of $d_n'(t)$ at $t = 0$ is:

$$(3.1.1) \qquad \sqrt{\frac{1}{\cosh^2(d_n/2)} - \frac{1}{\cosh^2 J(T_{i_n^+}(\mathbf{d}))}} - \sqrt{\frac{1}{\cosh^2(d_n/2)} - \frac{1}{\cosh^2 J(T_{i_n^-}(\mathbf{d}))}}$$

Above, $J(T_{i_n^+}(\mathbf{d}))$ is the circumcircle radius of $C_{i_n^+}$, and therefore also of $T_{i_n^+}$, and likewise for $J(T_{i_n^-}(\mathbf{d}))$. Lemma 2.3 of [3] implies that the former is greater than the latter, and it follows in this sub-case that the coefficient of $d'_n(t)$ is greater than 0 at $t = 0$.

It is also possible in this case that the Voronoi edge e geometrically dual to γ_n is noncompact. Then arguing as in the proof of [3, Lemma 5.8] establishes that $C_{i_n^+}$ is also non-compact: for the universal cover $\pi\colon \mathbb{H}^2 \to S$, if $\mathcal{S} = \pi^{-1}(x) \subset \mathbb{H}^2$ and $\tilde{e}$ is a lift of e, Lemma 1.9 of [3] asserts that its ideal endpoint v_∞ is the ideal point of a horocycle S with the property that the horoball B bounded by S satisfies $B \cap \mathcal{S} = S \cap \mathcal{S}$, and S contains the endpoints of the geometric dual γ to $\tilde{e}$. By the empty circumcircles condition, the convex hull of $B \cap \mathcal{S}$ is a non-compact Delaunay two-cell $\tilde{C}$ containing γ, which is a lift of γ_n since $\tilde{e}$ is a lift of e. Hence $\tilde{C}$ projects to a non-compact two-cell containing γ_n, necessarily $C_{i_n^+}$.

Corollary 2.1 implies that on the interior of $\tilde{C}$ the projection to $C_{i_n^+}$ factors through an embedding of $int(\tilde{C})/\Gamma$, where Γ is the stabilizer of v_∞ in $\pi_1 S$. We may assume that the triangulation of $C_{i_n^+}$ has been obtained by dividing $\tilde{C}$ into triangles with geodesic rays joining its vertices to v_∞, then projecting, so in particular $T_{i_n^+}$ is the projection of a horocyclic ideal triangle with compact side of length d_n. From the second equation of Proposition 3.7 of [4] we therefore obtain:

$$(3.1.2) \quad \frac{\partial}{\partial d_n}\left(D_0(T_{i_n^+}(\mathbf{d})) + D_0(T_{i_n^-}(\mathbf{d}))\right) = \frac{1}{\cosh(d_n/2)} - \sqrt{\frac{1}{\cosh^2(d_n/2)} - \frac{1}{\cosh^2 J(T_{i_n^-}(\mathbf{d}))}}$$

Again this is positive, and the claim follows in this case.

The second case of the claim is when neither of $C_{i_n^\pm}$ is non-centered with longest edge γ_n, whence the same holds for the $T_{i_n^\pm}$. In this case both terms of the coefficient of $d'_n(t)$ are positive, by Propositions 2.3 or 3.7 of [4], applied as above. $\qquad\square$

Proof of Theorem 0.1. Let (S, x) be a pointed surface whose Delaunay tessellation is not of the form described in the Theorem. The goal is to show that there is a deformation of (S, x) that increases injectivity radius at x. We will assume we are in the case not covered by Proposition 3.1: all Delaunay edges of (S, x) have length $2injrad_x(S)$, and there is a Delaunay two-cell C which is compact and not a triangle, or non-compact and not a monogon.

Let $\mathcal{T}$ be a triangulation that is compatible with the Delaunay tessellation of (S, x), with an edge γ_0 that is a diagonal of C with a compact triangle T_1 of $\mathcal{T}$ on one side and the remainder of C on the other. If C is non-compact we accomplish this as follows: for a locally isometric universal cover $\pi\colon \mathbb{H}^2 \to S$ and a horocyclic two-cell $\widetilde{C}$ of the Delaunay tessellation of $\pi^{-1}(x)$ projecting to C, let $\tilde{\gamma}_0$ join vertices of $\widetilde{C}$ separated by exactly one other vertex on the horocycle in which it is inscribed. Then the compact subregion $\widetilde{T}_1$ of $\widetilde{C}$ that it bounds is a triangle, hence so is its projection T_1. In this non-compact case we divide the remainder of C, and all other horocyclic Delaunay cells, into horocyclic ideal triangles as previously.

Enumerate the edges of $\mathcal{T}$ as $\gamma_0, \dots, \gamma_l$ so that the Delaunay edges are $\gamma_1, \dots, \gamma_n$ for some $n \le l$. Let d_j be the length of γ_j for each j, and note that by hypothesis γ_j has length $d \doteq 2injrad_x(S)$ for $0 < j \le n$. Now let $\mathbf{d} = (d_0, \dots, d_l)$, and prescribe $\mathbf{d}(t) = (d_0(t), \dots, d_l(t))$ with $\mathbf{d}(0) = \mathbf{d}$ as follows: $d_0(t) = d_0 - t$; $d_j(t) \equiv d_j$ for $j > n$; and for $0 < j \le n$, $d_j(t) = d(t)$ is determined by the differential equation $\frac{d}{dt} A_x(\mathbf{d}(t)) = 0$.

Here as in the proof of Proposition 3.1, A_x is the angle sum function from Definition 1.1, and for a smooth solution $\mathbf{d}(t)$ we have $A_x(\mathbf{d}(t)) \equiv 2\pi$ since $\mathbf{d}$ is the edge length collection of the triangulated hyperbolic surface $(S, \mathcal{T})$. It will then follow from Proposition 1.3 that $S(\mathbf{d}(t))$ is a deformation of S in $\mathfrak{T}_{g,n}$. And if $d(t)$ increases with t, then by Proposition 1.4, the injectivity radius of $S(\mathbf{d}(t))$ at its vertex will as well. We will show this below.

As in the proof of Proposition 3.1 we rewrite the equation $\frac{d}{dt} A_x(\mathbf{d}(t)) = 0$ using the chain rule and our choices for $\mathbf{d}(t)$:

(3.1.3)
$$
d'(t) \sum_{j=1}^{n} \frac{\partial}{\partial d_j} \left(D_0(T_{i_j^+}(\mathbf{d}(t))) + D_0(T_{i_j^-}(\mathbf{d}(t))) \right) - \frac{\partial}{\partial d_0} \left(D_0(T_0(\mathbf{d}(t))) + D_0(T_1(\mathbf{d}(t))) \right) = 0
$$

Again as in Proposition 3.1, for each $j > 0$ the $T_{i_j^\pm}$ are the triangles containing the edge γ_j. Here T_0 and T_1 are the triangles containing γ_0, and by construction, T_1 is compact. In all cases if $T_{i_j^\pm}$ has edges $\gamma_{j_1}, \gamma_{j_2}, \gamma_{j_3}$ then $T_{i_j^\pm}(\mathbf{d}(t))$ refers to the collection $(d_{j_1}(t), d_{j_2}(t), d_{j_3}(t))$ of changing edge lengths. We claim first that all coefficients above are smooth, and that the coefficient of $d'(t)$ is positive.

To the latter point, recall that since $\mathcal{T}$ is compatible with the Delaunay tessellation of (S, x), each $T_{i_j^\pm}$ is contained in a Delaunay cell $C_{i_j^\pm}$. If $C_{i_j^\pm}$ is compact it is centered, being equilateral, so since $T_{i_j^\pm}$ has the same circumcircle it is either centered or one of its edges is a diagonal that separates it from the circumcircle center. In neither of these cases is it non-centered with longest edge γ_j, so by Proposition 2.3 of [4] its contribution to the coefficient of $d'(t)$ is positive. If the Delaunay cell $C_{i_j^\pm}$ containing $T_{i_j^\pm}$ is horocyclic, and hence $T_{i_j^\pm}$ is a horocyclic ideal triangle by construction, then this follows from Proposition 3.7 of [4].

Smoothness of the coefficients of (3.1.3) follows from results of [4]. In particular, Proposition 2.3 there asserts that D_0 is smooth on the set $\mathcal{AC}_3 \subset (0, \infty)^3$ parametrizing cyclic triangles. This applies to each $T_{i_j^\pm}$ contained in a compact Delaunay cell. Each one contained in a horocyclic cell, except possibly T_1, is a horocyclic ideal triangle by construction, and smoothness follows by [4, Prop. 3.7]. If T_1 is in a horocyclic Delaunay cell then its side-length collection $T_1(\mathbf{d})$ lies in the set $\mathcal{HC}_3$ of [4, Corollary 3.5], parametrizing compact "horocyclic" triangles. $\mathcal{HC}_3$ has codimension one in $(0, \infty)^3$. It bounds the set $\mathcal{C}_3$ parametrizing cyclic triangles on one side, and the set $\mathcal{E}_3$ parametrizing "equidistant" triangles on the other (see [4, §4]; in particular Cor. 4.6 there).

For arbitrary $n \geq 3$, the versions of D_0 that record areas of horocyclic and equidistant n-gons are respectively defined in Propositions 3.7 and 4.9 of [4]. We proved there that the various definitions of D_0 determine a continuous function on $\mathcal{AC}_n \cup \mathcal{HC}_n \cup \mathcal{E}_n$, but we did not address smoothness on $\mathcal{HC}_n$. However since D_0 measures area, for $n = 3$ it agrees everywhere with the smooth function A of [4, Lemma 1.16]. Therefore since T_1 is a triangle, the coefficient function $D_0(T_1(\mathbf{d}(t)))$ of (3.1.3) is smooth.

Since the coefficient of $d'(t)$ in (3.1.3) is positive at $\mathbf{d}$ and all coefficients are smooth, there is a smooth solution $d(t)$ near $t = 0$. The sign of $d'(t)$ is determined by the sign of $\frac{\partial}{\partial d_0} \left(D_0(T_0(\mathbf{d}(t)) + D_0(T_1(\mathbf{d}(t))) \right)$.

Claim 3.1.4. For small $t > 0$, $\frac{\partial}{\partial d_0} \left(D_0(T_0(\mathbf{d}(t)) + D_0(T_1(\mathbf{d}(t))) \right) > 0$, hence $d'(t) > 0$.

Proof of claim. First suppose C is non-compact. Then $d'(0) = 0$, since at time 0 we have:

$$\frac{\partial}{\partial d_0} D_0(T_0(\mathbf{d})) = \frac{1}{\cosh(d_0/2)} = -\frac{\partial}{\partial d_0} D_0(T_1(\mathbf{d}));$$

The computation here for T_0 is obtained by taking a derivative with respect to d_0 of the second formula of [4, Proposition 3.7]. For T_1 it follows similarly from the first formula there, noting that by construction d_0 is the largest side length of T_1.

This requires some comment since the formula in question applies only to points of $\mathcal{HC}_n$, which as we pointed out above is codimension-one in $(0, \infty)^n$. But since we have chosen γ_0 so that $n = 3$, as pointed out above D_0 is smooth on a neighborhood of the side length collection $T_1(\mathbf{d}) = (d_0, d, d)$ of T_1, and its partial derivative with respect to d_0 at this point is a limit of $\frac{\partial D_0}{\partial d_0}(\mathbf{d}_n)$ for a sequence $\{\mathbf{d}_n\} \in \mathcal{AC}_3$ approaching $T_1(\mathbf{d})$. Noting that all but finitely many $\mathbf{d}_n$ are in $\mathcal{AC}_3 - \mathcal{C}_3$, by Corollary 3.5 of [4], and the circumcircle radius $J(\mathbf{d}_n) \to \infty$ as $n \to \infty$, by Proposition 3.6 there, the given formula is a limit of the one given by Proposition 2.3 there.

By the above we have that $\frac{d}{dt}T_1(\mathbf{d}(0)) = (-1, 0, 0)$. Near (d_0, d, d), $\mathcal{HC}_3$ is characterized as a graph $\{h_0(x, y), x, y)\}$ by Corollary 3.5 of [4], for h_0 as defined there, and $\mathcal{AC}_3$ is characterized as $\{(x, y, z) \,|\, x < h_0(y, z)\}$; compare with [4, Corollary 1.10]. Thus this vector points into $\mathcal{AC}_3$, so $T_1(\mathbf{d}(t)) \in \mathcal{AC}_3$ for all small-enough $t > 0$. For all such t it follows that $\frac{\partial}{\partial d_0}(D_0(T_0(\mathbf{d}(t)) + D_0(T_1(\mathbf{d}(t))))$ is given by the formula of (3.1.2), with d_0 replacing d_n and $T_0(\mathbf{d}(t))$ replacing $T_{i_n^-}(\mathbf{d}_0)$ there. This quantity is positive, therefore so is $d'(t)$, and the claim holds if C is non-compact.

We now address the case that C is compact. First suppose that C is a quadrilateral. By hypothesis all its edge lengths are equal to $d = 2injrad_x(S)$, so since it is cyclic and therefore uniquely determined by its edge length collection it is fully symmetric. In particular, each diagonal of C is a diameter of its circumcircle, so $d_0 = 2J(d_0, d, d)$, where $J \colon \mathcal{AC}_n \to (0, \infty)$ records circumcircle radius of cyclic polygons; see Proposition 1.14 of [4]. Plugging this into Proposition 2.3 there gives $d'(0) = 0$ again.

Again in this case we have $\frac{d}{dt}T_1(\mathbf{d}) = (-1, 0, 0) = \frac{d}{dt}T_0(\mathbf{d})$. In this case the edge length collections of T_0 and T_1 lie in the set $\mathcal{BC}_3$ parametrizing *semicyclic* triangles, cyclic triangles with one side a diameter of their circumcircles. This is a codimension-one submanifold of $(0, \infty)^3$ which is the frontier of $\mathcal{C}_3$, the open set parametrizing centered triangles, in $\mathcal{AC}_3$; see [4, Proposition 1.12]. The vector $(-1, 0, 0)$ points into $\mathcal{C}_3$ at $T_1(\mathbf{d})$, since near here $\mathcal{BC}_3$ is a graph $\{(b_0(x, y), x, y)\}$ (see [4, Prop. 1.12]) and $\mathcal{C}_3 = \{(z, x, y) \,|\, z < b_0(x, y)\}$ (compare [4, Prop. 1.11]). Therefore $T_1(\mathbf{d}(t)) = T_0(\mathbf{d}(t)) \in \mathcal{C}_3$ for all small $t > 0$, and it follows from Proposition 2.3 of [4] that $d'(t) > 0$ for such t.

If C is not a quadrilateral then we may choose γ_0 and T_1 so that the circumcircle center of C lies on the opposite side of γ_0 from T_1. Then $T_1(\mathbf{d}) \in \mathcal{AC}_3 - (\mathcal{C}_3 \cup \mathcal{BC}_3)$ has largest entry d_0. On the other hand either $T_0(\mathbf{d}) \in \mathcal{C}_3$, i.e. T_0 is centered, or T_0 is not centered and γ_0 is not its longest side. The condition on $T_1(\mathbf{d})$, being open, holds for $T_1(\mathbf{d}(t))$ for all t near 0. Similarly, if $T_0(\mathbf{d}) \in \mathcal{C}_3$ then this also holds for $T_0(\mathbf{d}(t))$, or if d_0 is not the largest entry of $T_0(\mathbf{d})$ then $d_0(t)$ is not the largest entry of $T_0(\mathbf{d}(t))$, for all t near 0. Proposition 2.3 of [4] thus implies that $\frac{\partial}{\partial d_0}(D_0(T_0(\mathbf{d}(t)) + D_0(T_1(\mathbf{d}(t))))$ is given for all such t by the formula (3.1.1), with d_0 replacing d_n, $T_0(\mathbf{d}(t))$ replacing $T_{i_n^+}(\mathbf{d})$, and $T_1(\mathbf{d}(t))$ replacing $T_{i_n^-}(\mathbf{d})$.

We have $d'(0) = 0$ by (3.1.1), since T_0 and T_1 are both inscribed in the circumcircle of C. For $t \neq 0$, if $J(T_0(\mathbf{d}(t))) > J(T_1(\mathbf{d}(t)))$ then $d'(t) > 0$, again by (3.1.1). Applying Proposition 1.14 of [4], we obtain either

$$0 < \frac{\partial}{\partial d_0} J(T_0(\mathbf{d})) < 1/2 < \frac{\partial}{\partial d_0} J(T_1(\mathbf{d})),$$

if T_0 is centered (i.e. $T_0(\mathbf{d}) \in \mathcal{C}_3$), or $\frac{\partial}{\partial d_0} J(T_0(\mathbf{d})) < 0$ if not. Since $d'(0) = 0$, the chain rule implies that $\frac{d}{dt} J(T_0(\mathbf{d})) = -\frac{\partial}{\partial d_0} J(T_0(\mathbf{d}))$, and similarly for $\frac{d}{dt} J(T_1(\mathbf{d}))$. Thus for $t > 0$, $J(T_0(\mathbf{d})) > J(T_1(\mathbf{d}))$ so $d'(t) > 0$, and the claim is proved in all cases. $\square$

Lemma 1.2 now implies that $\mathbf{d}(t)$ determines a path $(S(\mathbf{d}(t)), \mathcal{T}(\mathbf{d}(t)))$ of triangulated hyperbolic surfaces, which Proposition 1.3 implies is continuous in $\mathfrak{T}_{g,n}$. By Proposition 1.4 and our construction, S_t has injectivity radius $d(t)/2$ at the vertex of $\mathcal{T}_t$, so since d increases with t the result holds. $\square$

References

[1] Christophe Bavard. Disques extrémaux et surfaces modulaires. *Ann. Fac. Sci. Toulouse Math. (6)*, 5(2):191–202, 1996.

[2] Jason DeBlois. The Delaunay tessellation in hyperbolic space. Preprint. arXiv:1308.4899.

[3] Jason DeBlois. The centered dual and the maximal injectivity radius of hyperbolic surfaces. *Geom. Topol.*, 19(2):953–1014, 2015.

[4] Jason DeBlois. The geometry of cyclic hyperbolic polygons. *Rocky Mountain J. Math.*, 46(3):801–862, 2016.

[5] Jason DeBlois. Bounds for several-disk packings of hyperbolic surfaces. Preprint. arXiv:1701.07770, January 2017.

[6] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.

[7] Matthieu Gendulphe. The injectivity radius of hyperbolic surfaces and some Morse functions over moduli spaces. Preprint. arXiv:1510.02581, October 2015.

[8] Svetlana Katok. *Fuchsian groups*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992.

[9] John G. Ratcliffe. *Foundations of hyperbolic manifolds*, volume 149 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.

[10] P. Schmutz. Riemann surfaces with shortest geodesic of maximal length. *Geom. Funct. Anal.*, 3(6):564–631, 1993.

[11] William P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1997. Edited by Silvio Levy.

Department of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260

E-mail address: jdeblois@pitt.edu

URL: http://www.pitt.edu/~jdeblois